

1 **CLUSTER-BASED CACHE MEMORY ALLOCATION**

2 This application is a continuation of Application No. 09/522,407, filed April 19, 2000,
3 herein incorporated by reference.

4 **BACKGROUND OF THE INVENTION**

5 **Field of the Invention**

6 The present invention relates to disk drive performance features, and more particularly to
7 a disk drive having a cache control system for improving the disk drive's response time to host
8 commands.

9 **Description of the Prior Art**

10 A host computer stores and accesses data on a disk drive by issuing commands
11 to the disk drive over a standardized interface. The smallest indivisible data unit addressable on
12 a disk is a logical block or disk sector, typically of 512 bytes, and each such disk sector is
13 assigned a logical block address (LBA). When the host computer sends a command to the disk
14 drive, the nature of the command is specified, e.g., read or write, along with a start LBA and a
15 count specifying the number of contiguous sectors to be transferred.

16 Existing disk drives typically have a semiconductor cache memory for temporarily storing
17 disk data that is likely to be requested by a host computer. The response time latency for storing
18 and accessing data in a semiconductor memory is much smaller than the response time latency
19 for mechanically storing and accessing data stored on a rotating disk. In existing disk drives,
20 disk data is generally cached in contiguous fixed length memory segments. The memory
21 segments may be inefficiently configured to order to accommodate host commands having a long
22 LBA range thereby wasting valuable data storage space in the cache memory if such commands
23 occur infrequently.

24 Accordingly, there exists a need for a disk drive having a disk cache system for efficiently
25 allocating and configuring memory segments for responding to host commands. The present
26 invention satisfies these needs.

SUMMARY OF THE INVENTION

The invention may be embodied in a disk drive, and related method, for servicing host disk commands using a cache memory having a plurality of sequentially-ordered memory clusters for caching disk data of disk sectors identified by logical block addresses. The disk drive includes a cache control system having a plurality of cluster control blocks and a tag memory usable only for providing a plurality of tag records. Each cluster control block has a cluster segment record for associating the cluster control block with a particular memory cluster and for forming variable length segments of the memory clusters without regard to the sequential order of the memory clusters. Each tag record assigns a segment to a contiguous range of logical block addresses and defines the cluster control blocks forming the segment. Each segment of the memory clusters is for caching disk data of the assigned contiguous range of logical block addresses.

In more detailed features of the invention, the cluster segment record of each cluster control block associated with a segment may include a pointer to a subsequent cluster control block or to indicate an end cluster control block of the segment. Each tag record may define a length for the assigned segment by pointing to a first cluster control block and by pointing to a last cluster control block for the segment or by indicating a count of cluster control blocks for the segment. Also, the memory clusters may be uniformly sized.

In other more detailed features of the invention, the disk drive may further include a free list for identifying cluster control blocks not associated with a tag memory record and forming a segment of the identified cluster control blocks. Each tag record for an assigned segment may point to a first cluster control block, a last cluster control block, and indicate a count of sectors for the segment. A length of an original assigned segment may be increased by removing a segment of cluster control blocks from the free list segment, changing the original segment's last cluster control block to point to a first cluster control block of the removed segment, and changing the tag record to point to the last cluster control block of the removed segment. The disk drive also may include a microprocessor that de-allocates an existing assigned segment and assigns the segment's associated cluster control blocks to the free list if a sufficient number of cluster control blocks are not available on the free list to enable caching of a range of logical block addresses requested by a host command.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings illustrate embodiments of the present invention and, together with the description, serve to explain the principles of the invention.

FIG. 1 is a block diagram of a disk drive having a cache control system with cluster control blocks (CCBs) for forming variable-length memory segments using memory clusters without regard to the sequential order of the memory clusters, according to the present invention.

FIG. 2 is a block diagram showing the cache control system of FIG. 1, with a tag memory having tag records and a CCB memory having the CCBs, according to the present invention.

FIG. 3 is a block diagram showing the relationship between the CCBs and the memory clusters.

FIG. 4 is a data structure for a CCB for forming memory segments.

FIG. 5 is a block diagram of a free list for forming a segment of CCBs not assigned to a tag record.

FIG. 6 is a block diagram showing a table of tag records in a tag memory of the cache control system of FIG. 1, for defining the CCBs forming a segment of memory clusters.

FIG. 7 is a data structure for a tag record in the table of tag records of FIG. 5, for forming memory segments using the CCBs.

FIG. 8 is a flow chart showing a method for allocating CCBs to form variable length memory segments without regard to the sequential order of the memory clusters in the segments.

DETAILED DESCRIPTION

With reference to FIG. 1, a disk drive 10 comprises a cache memory 14 having a plurality of sequentially-ordered memory clusters 46 for caching disk data stored in sectors (not shown) on disks of a disk assembly 38. Conventionally the disk sectors are identified by logical block addresses (LBAs). A cache control system 12 of the disk drive 10 comprises a cluster control block memory, having plurality of cluster control blocks (CCBs) 34, and a tag memory 22, having a plurality of tag records 40, that are embedded within the cache control system 12. Each CCB 34 includes a cluster segment record with an entry 90 (FIG. 4) for associating the CCB 34 with a particular memory cluster 46 and for forming variable length segments of the memory clusters 46 without regard to the sequential order of the memory clusters 46. Each tag record 40 assigns a segment to a continuous range of LBAs and defines the CCBs 34 forming the segment.

1 Each segment of the memory clusters 46 is for caching data from a contiguous range of the
2 logical block addresses. The cache control system 12 efficiently exploits available memory
3 clusters 46 for responding to host commands.

4 The disclosures of the following three U.S. Patent Applications are hereby incorporated
5 by reference: application serial number 09/552,404, filed April 19, 2000, now patent number
6 6,553,457, titled TAG MEMORY DISK CACHE ARCHITECTURE; application serial number
7 09/552,399, filed April 19, 2000, now patent number 6,601,137, titled RANGE-BASED CACHE
8 CONTROL SYSTEM AND METHOD; and application serial number 09/552,402, filed
9 April 19, 2000, titled CACHE CONTROL SYSTEM AND METHOD HAVING HARDWARE-
10 BASED TAG RECORD ALLOCATION.

11 With reference again to FIG. 1, the disk drive 10 further includes a microprocessor 16,
12 and a host interface 18. The host interface 18 receives host commands from a host 20, such as a
13 personal computer, and transfers disk data between the disk drive 10 and the host 20. The host
14 commands identify the disk data using a start logical block address (LBA) and a count specifying
15 the number of contiguous sectors to be transferred. The cache memory 14 caches the disk data
16 under the direction of the cache control system 12 and the microprocessor 16. The
17 microprocessor 16 operates under firmware control and manages the operation of the disk drive
18 10 and assists hardware elements under specific conditions. The cache memory 14 is random
19 access memory, typically 2 megabytes (MB). Generally, the larger the cache memory 14, the
20 better the performance of the disk drive 10 in responding to host commands.

21 The disk drive 10 also includes a disk channel 36 and the aforementioned disk assembly
22 38. The disk assembly 38 includes a disk platter that is organized into the disk sectors, typically
23 of 512 bytes plus redundancy bytes for error correction, which are individually addressable using
24 a logical block address (LBA). The disk channel 36 performs conventional encoding and
25 decoding of data written to and read from the disk.

26 The cache control system 12 is shown in more detail in FIG. 2. The cache control system
27 12 includes the aforementioned tag (random access) memory (RAM) 22 and the aforementioned
28 CCB memory or RAM 24. The tag memory 22 is a static random access memory (SRAM)
29 structure which is preferably embedded in an integrated controller chip having a table of the
30 aforementioned tag or segment records 40. The CCB memory is also preferably an embedded

1 SRAM. The embedded tag memory 22 and CCB memory 24 thus provides higher performance
2 and lower cost versus firmware based cache control schemes which use a general-purpose
3 external RAM. In particular, since internal hardware engines, as described further below, may
4 access the tag and CCB records independently from microprocessor 16, the cache control system
5 12 enables higher performance by off-loading microprocessor 16 and providing hardware-based
6 processing as detailed below.

7 As shown in FIG. 3, the cache memory 14 is divided into 512 byte blocks or sectors 48.
8 Each cache sector 48 is for storing disk data of a 512 byte disk sector. Note that if the number of
9 bytes in a disk sector is defined to have, for example, 1024 bytes, then the number of bytes in a
10 cache sector is similarly defined to have 1024 bytes. The cache sectors 48 are bunched into the
11 aforementioned consecutively numbered groups or clusters 46. Each cluster 46 has a particular
12 cluster number. Preferably, each cluster 46 has 16 cache sectors 48, although the number of
13 sectors 48 in each cluster 46 may be selected based on the size of the cache memory 14, the size
14 of the CCB SRAM 24, and the operational characteristics of the host 20. The CCB memory 24 is
15 initialized by the microprocessor by forming records or entries for the CCBs 34, shown in FIG. 4.
16 Each CCB record indicates a cluster number 90 and includes a next CCB pointer 62. Other CCB
17 record entries are included for indicating the number of valid sectors 92 and a skip count 94 to
18 the first valid sector 48 in the CCB, and for pointing to a next sector to be piped 96 and a
19 remaining count 98 of sectors 48 to be piped, for data transfer operations.

20 The cache control system 12 also includes a free list 64 shown in FIG. 5. The free list 64
21 tracks any CCBs 34 not assigned to a tag record 40. The CCB memory 24 is initialized by the
22 microprocessor 16 with all of the CCBs 34 being initially assigned to the free list 64. The free
23 list 64 includes entries for a CCB count 84, a first CCB number 86 and a last CCB number 88.
24 The count entry 84 indicates the number of CCBs 34 on the free list 64 available for assignment
25 to tag records 40. The first CCB entry 86 points to the first CCB 34 on the free list 64. The last
26 CCB entry 88 points to the last CCB 34 on the free list 64. Each CCB 34 points to the next CCB
27 34 in the free list 64 for forming a free list segment 82. At initialization, the CCBs 34 in the free
28 list segment 82 are configured in consecutive order, but the order of the CCBs in the free list 64
29 may be scattered after numerous CCB 34 assignments to and returns from the varying length
30 segments defined by the tag records 40. All CCBs 34 are assigned to either a tag record 40 or to

1 the free list 64.

2 In the exemplary free list segment 82 shown in FIG. 5, the free list's first CCB 86 entry
3 points to the CCB number 45. The CCB number 45 points to the CCB number 6. The free list
4 segment 82 continues until the last CCB number 95. The CCB number 95 has a null value in its
5 next CCB pointer 62. The free list's last CCB entry 88 points to the CCB number 95. The length
6 of a memory segment formed by the CCB's 34 may be increased by removing CCB's 34 from the
7 free list segment 82, changing the last CCB entry 88 to point to the new last CCB, updating the
8 count entry 84 by subtracting the number of CCBs removed from the free list segment 82, and
9 changing the next segment pointer 62 of the new last CCB to the null value.

10 The cache control system 12 also includes a scan engine 26 and a host writable control
11 store (HWCS) 28. The scan engine 26 is coupled to the host interface 18 and receives host
12 commands and scans the tag memory 22 for the LBA ranges associated with a host command.
13 The scan engine 26 places the scan results in a results register 30 or, if servicing the host
14 command further requires intervention by the microprocessor 16, the HWCS 28 places the
15 command in a command queue 32. The command queue 32 has a read miss queue and a write
16 command first-in first-out (FIFO) queue. The scan engine 26 is described in more detail in the
17 above-referenced U.S. patent number 6,601,137, titled RANGE-BASED CACHE CONTROL
18 SYSTEM AND METHOD. The CCB memory 24 may be updated by the microprocessor 16 and
19 by the HWCS 28. If a tag record 40 may be allocated for responding to a host command, then
20 the HWCS 28 manages the response to the host command, otherwise the microprocessor 16 may
21 assist with the response. Thus, the HWCS 28 off-loads cache tasks from the microprocessor 16
22 enabling response to host commands for data without microprocessor intervention.

23 The tag memory 22 is described in more detail with reference to FIGS. 6 and 7. The tag
24 memory 22 has the plurality of tag records 40 that define segments, 42 and 44, of memory
25 clusters 46 within the cache memory 14. Typically, the tag memory 22 may have 32 or 64
26 records dedicated to defining variable length segments. Other tag memory records (not shown)
27 may be dedicated to single block transfers for caching small data elements stored within one
28 memory cluster 46 that are repeatedly accessed by the host 20.

29 The tag memory 22 defines the segments of the cache memory clusters 46 using the
30 CCBs 34. Each tag record 40 has entries or fields (50, 52, 54, 56, 58 and 60), respectively for

1 indicating the first disk LBA assigned to the corresponding segment, the number of valid sectors
2 in the segment, the number of sectors allocated to the segment, the first segment CCB, the last
3 segment CCB, and state and control flags for the segment. A tag record 40 defines a segment by
4 recording the segment's first CCB in the first CCB entry 56. The first CCB 34 has a pointer 62
5 to the next or second CCB in the segment. The second CCB likewise has a pointer 62 to the next
6 CCB until the last CCB in the segment. The last CCB has an indicator such as a null value that
7 indicates the end of the segment.

8 Two short exemplary segments, 42 and 44, are shown in FIG. 3. The first segment 42 is
9 defined by the tag record number 1 to have a length of three clusters 46. The second segment 44
10 is defined by the tag record number 29 to have a length of two clusters 46. The last CCB in each
11 segment, 42 and 44, has a null value in the next CCB pointer 62. The tag memory 22 is
12 described in more detail in the above-referenced U.S. patent number 6,553,457 titled TAG
13 MEMORY DISK CACHE ARCHITECTURE.

14 The length of a segment defined by a tag record 40 may be extended by changing the last
15 CCB 34 of the segment to point to a next CCB added from the free list 64 and by updating the
16 last segment CCB entry 58 in the tag record 40. The microprocessor 16 may set a threshold
17 count for the CCBs 34 on the free list 64. If the free list count entry 84 falls below the threshold,
18 the microprocessor 16 may de-allocate tag records 40 thus freeing up CCBs 34 to avoid an
19 insufficient CCB error.

20 The tag record pointer entries, 56 and 58, in conjunction with the CCB pointers 62, allow
21 definition of variable length segments without regard to the logical or numerical order of the
22 clusters 46 in the cache memory 14. Accordingly, the tag memory 22 and the CCB memory 24
23 provide a flexible and powerful disk cache technique for efficiently responding to host
24 commands.

25 The cache control system 12 also includes a most-recently-used/least-recently-used
26 (MRU/LRU) engine 66. The MRU/LRU engine 66 keeps track of the usage of the cached data
27 associated with each tag record 40 in the tag memory and is described in detail in the above-
28 reference U.S. application serial number 09/552,402, titled CACHE CONTROL SYSTEM AND
29 METHOD HAVING HARDWARE-BASED TAG RECORD ALLOCATION.

30 As shown in FIG. 8, the invention also may be embodied in a method for caching disk

1 data of disk sectors identified by logical block addresses using a cache memory 14 having a
2 plurality of sequentially-ordered memory clusters 46. A plurality of cluster control blocks 34 are
3 provided (step 172) for forming memory cluster segments. Each cluster control block 34 has a
4 cluster segment record 90 for associating the cluster control block 34 with a particular memory
5 cluster 46 and for forming variable length segments of the memory clusters 46 without regard to
6 the sequential order of the memory clusters 46. A segment is assigned to a contiguous range of
7 logical block addresses (step 174). The cluster control blocks 34 forming the segment are
8 defined (step 176). Each segment of the memory clusters is for caching disk data of the assigned
9 contiguous range of logical block addresses.